CHAPTER 1

Working in 3D

The initial impulse for someone interested in learning 3D graphics is to immediately jump in and start building—trust me, I have been there! While there is definitely something good to be said about diving right in, if you spend a little time creating a foundation of skills, you will have less need to break bad habits later. Before you start creating anything, you need to understand the canvas you will be working with and the elements that will be used in the creative process.

This chapter covers the following:

- **Understanding 3D space**
- **Exploring a 3D scene**
- **Navigating the modo user interface and its viewports**
- **Maneuvering views and objects in space**

COPYRIGHTED MATERIAL

Understanding 3D Space

Maneuvering in 3D space can seem easy at first glance. After all, only one dimension has been added to the standard page layout, and that dimension is what we experience as we move around every day. However, believe it or not, the addition of this dimension can make navigation harder to get used to for the novice 3D artist. If you don't grasp some foundational principles from the outset, you can become disoriented and lose track of your model and scene. The addition of a third dimension adds much more than just another arrow on the monitor.

A standard page layout has two axes: x and y. If these equate to the horizontal and vertical directions, respectively, then the third axis (z) extends off the screen, toward the viewer (see Figure 1.1).

This works well for starters, but let's take it a step further and look at space in terms of a map or other top-down design. In this case, the plane defined by the x and z axes makes up the Cartesian plane. More specifically, the negative z-axis is north, and the negative x-axis is west, in relation to the middle of our workspace (see Figure 1.2).

The center of space (called the *origin*) will be the starting point for all of our design work unless we specifically need to work in a different area of the scene. Even in this case, it is often best to create an object at the origin and then move it to the desired location, because this will enable you to work with symmetry, easily locate objects, and move the mesh layer, which can then be reset or animated much more easily.

Figure 1.1
Three dimensions from a screen perspective

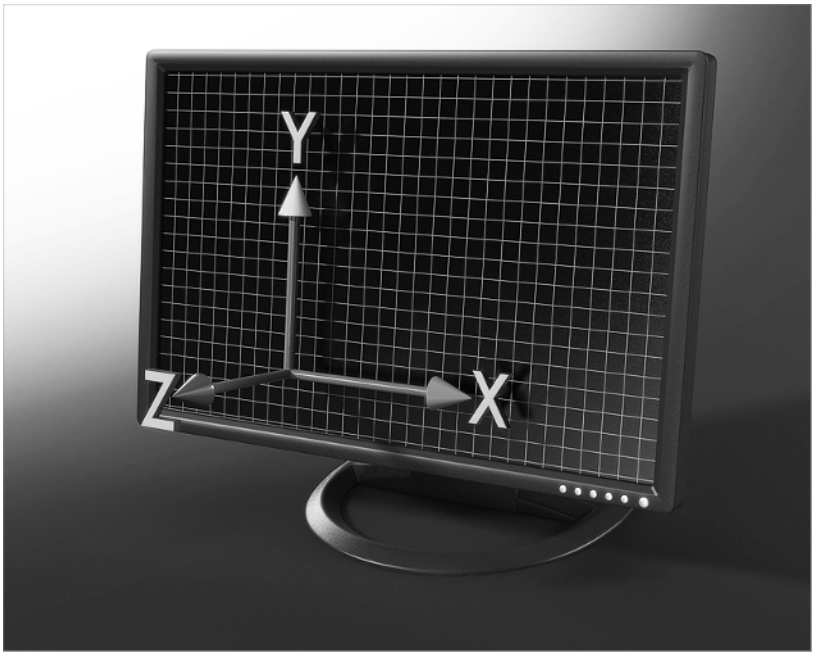

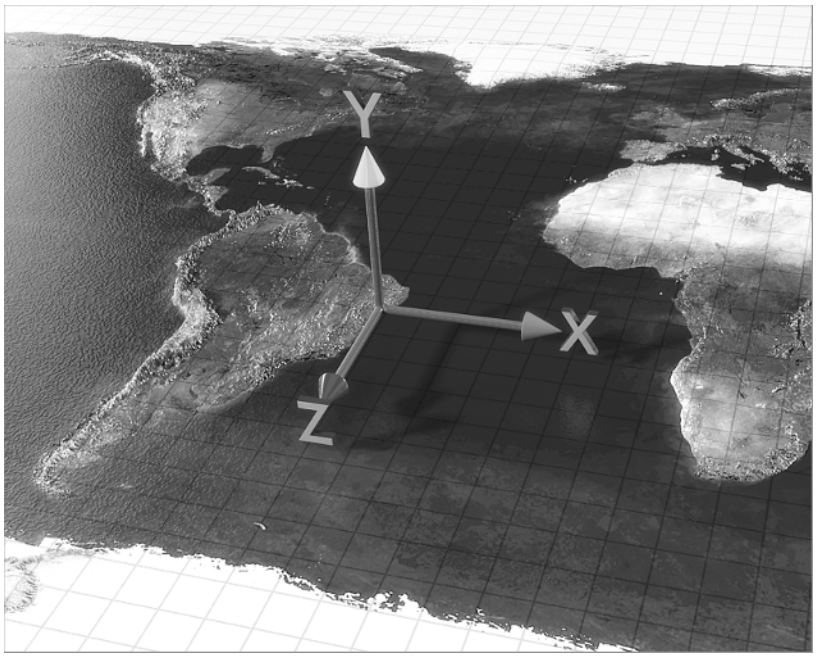

Figure 1.2
Three dimensions from a map perspective

Luxology's modo and many other common 3D applications work by default with the y-axis pointing up. Some applications, however, use z as the up-axis. The modo program allows you to customize this aspect in its preferences. Choose System → Preferences → Input → Accuracy And Units. From the heading marked Coordinate System, you can change the up-axis to z, y, or even x to suit your needs.

You can view a 3D scene through either an orthographic or a perspective view. An *orthographic view* offers a completely flat vantage point of objects, and placement without perspective of any kind. This means that objects located farther from the viewer will not appear smaller as the distance increases. An orthographic view is similar to a floor plan or elevation in architecture. Because it lacks the distortion associated with perspective, this type of view is ideal for creating and aligning objects.

Modo offers different interface layouts under its viewport tabs. Model Quad gives us three orthographic views and one perspective view. Although these two-dimensional views are initially set to Top, Front, and Right views, they can be changed to any other angle (Bottom, Left, or Back) and to views that include three-dimensional perspective.

Perspective views enable you to see objects and scenes with real depth. There are options to use an arbitrary perspective (the default in the Model Quad layout), camera perspective (based on the default scene camera or any additional cameras that have been added in the creative process), or light perspective. You can completely adjust the first

option without changing the scene in any way, but it is important to note that both camera and light perspectives are tied to actual objects in 3D space. Thus, if you move these views, you will actually be moving (or rotating) objects and changing the makeup of the scene. Movement in camera and light perspectives can be reversed with the Undo command, but movement in the generic perspective view cannot be undone.

A common analogy is that if an Adobe Photoshop (or other 2D graphic) image is like a painting, a 3D document is like a sculpture. Two-dimensional art forms (digital and analog) use space, form, and color to create the finished image. Depth and dimension are created through color variation for simulation of light and shadow. Three-dimensional art adds volume to the mix, which offers more-concrete simulation of real light and shadow. Because 3D provides added levels of realism, you need to consider additional parts of the creative process in order to create compelling 3D art.

Exploring a 3D Scene

The next things to consider when beginning to work in 3D are the individual pieces of a 3D scene and the steps in the creative process that will result in a finished project (a model, still image, animation, and so forth). This space is filled with points, edges, and polygons that create the objects. Materials and textures control the appearance of objects. Lights add shading and highlights to the scene. The camera provides the vantage point for the finished scene. Let's look at each of these in more depth.

Points, Edges, and Polygons

At the heart of any model that you will create are three basic elements: points (also called *vertices*), edges, and polygons. Points represent a single location in space. The initial impulse is often to think of these as being analogous to pixels in a raster image, but this is not the case. Because a vertex represents a single point in space, it is infinitesimally small and therefore does not appear in a finished 3D render. The fact that the vertex is the basic building block of the 3D creative process means that it is visible only when several are combined to create edges and polygons.

Edges make up the next level in the 3D food chain. When two points are joined together, an *edge* is created. This edge now exists in one dimension and is still invisible to the finished rendering. Once again, this should not be confused with a line in a 2D image file. For both edges and points, it is better to consider a vector illustration without any line weight assigned. Individual points and lines may appear on the page, but they will not print unless some thickness is attached to these elements. To see these elements, you must have a combination of at least three of them (points and edges), which creates a polygon.

A simple triangle represents the *polygon* at its most basic level. Three points (with three edges connecting them) creates a defined surface. For the most part, four-sided polygons, also known as quads, will be the basis for your models. The reason for this will become clear when modeling is discussed in the coming chapters, but suffice it to say that many forms can be more easily defined by quads than by triangles. By combining and blending together multiple polygons, objects take form.

SINGLE-SIDED POLYGONS

Polygons exist as two-dimensional elements within the three dimensions of a scene. The flat surfaces of polygons face in a single direction. Just as the points are infinitesimally small, so polygons are infinitely thin. This means that they are invisible when viewed from the back. Some thickness must be added in order to make the geometry appear from all angles.

Materials and Textures

After polygons are created, they must be assigned surface attributes to define their appearance. A *material* contains the basic description of how light interacts with a surface. The key components of a material are color, reflection, transparency, refraction, absorption, and emission of light. A material creates these attributes at a very basic level that is defined by either a color or a percentage (depending on the attribute). Proper combination of these properties can create a wide variety of looks and styles. To achieve something beyond the evenly distributed appearance of a basic material, additional layers must be added.

Textures add additional detail to surfaces. Textures are made from either rasterized 2D images or mathematical functions that display colors based on various inputs. Images can be placed on the surface of 3D models and offer a high degree of customization. You can place details exactly where you want them and edit them either by using an application such as Photoshop or by using texture painting inside of modo. The downside of image textures is that they can become pixilated if they are not of a high-enough resolution. Mathematical textures (known as procedural textures) are free from resolution and have a fairly wide range of styles, from simple grids and gradients to complex fractal algorithms. These textures, however, cannot be edited directly, so placement of detail is random.

These textures can be used to modulate any aspect of a material. Color can be applied to add variation as well as to colorize reflections or transparent tints. Other possibilities include changing the amount of reflection or transparency, the shininess, the translucency, or even adding the appearance of depth on a surface (see Figure 1.3).

Figure 1.3
A few examples of possible textures

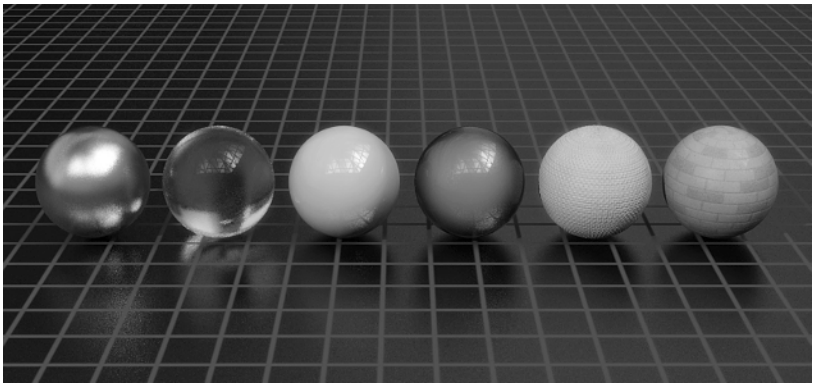

Lights

There are two methods of adding light to your 3D scenes: using computer representations of real lights (*standard lighting*) or casting light from the environment and textures in the scene (radiosity). The former is relatively easy to compute and delivers results more quickly. The latter uses more-complex computation and slows the finished image but results in lighting with more subtlety, nuance, and realism. Figure 1.4 shows a simple scene with standard lighting, and Figure 1.5 shows that same scene using radiosity.

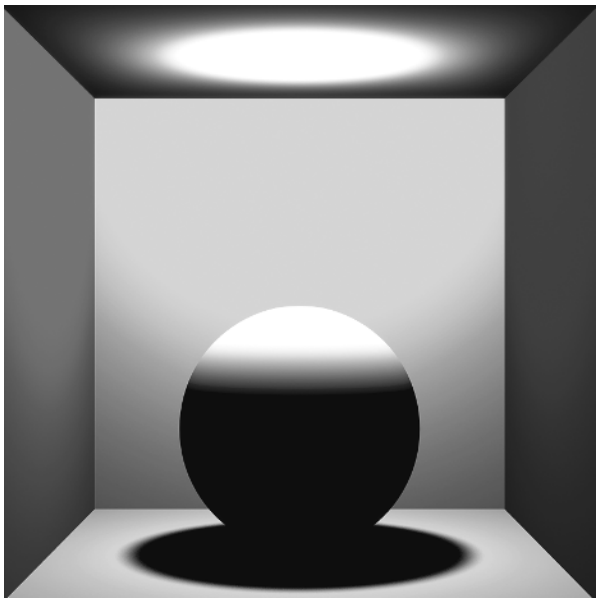

Figure 1.4
A scene using standard 3D lights

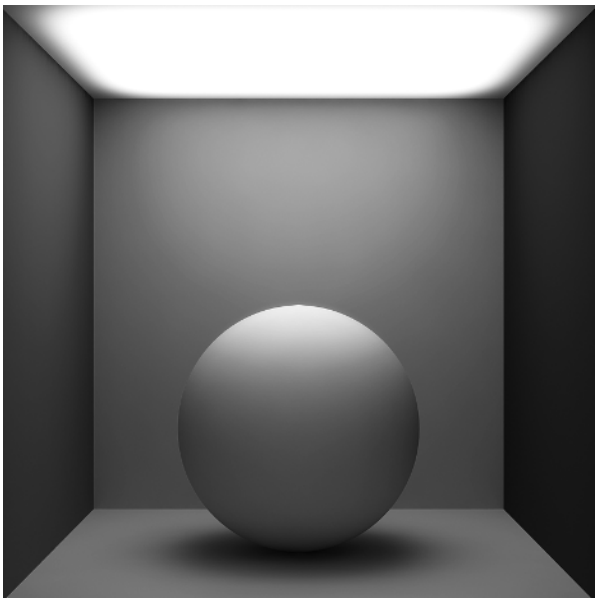

Figure 1.5
A scene using radiosity for lighting

Traditional lights in 3D space use simple math to add brightness based on an area of influence, light color, and intensity. These lights come in common variations that are seen in most 3D applications. *Distant lights* (sometimes called directional lights) are similar to the sun. The actual light comes from infinitely far away and is adjusted by the angle it enters the scene. *Spotlights* simulate their real-world namesake. Point lights are similar to lightbulbs and cast light outward from a single point, in all directions evenly. *Area lights* are similar to soft boxes used in studios. Other lights are more situational and are covered in depth (along with the other lights mentioned here) in the following chapters.

Radiosity is also known as global illumination, because illumination comes from other angles than just the direct light source. In this lighting model, light is based on light particles (known as *samples*), which project into a scene much like real-world light. As with light particles in the real world, these samples can bounce off surfaces to provide illumination in areas where a light does not have a direct effect. Each bounce of light adds an order of complexity to the calculation and, as a result, causes a slower render time. Because light samples are blended together for the final result, using more of them creates a smoother finished look and (like other quality-improving options) slows the final image production.

In general, a combination of both lighting types gives the best quality and control. However, there are times when using just one or the other can deliver excellent results.

The Camera

The camera in a 3D scene gives the viewpoint for finished images. Cameras appear only as representations in the scene and will not appear when a finished image is rendered (so you don't have to worry about them showing up in reflections). Cameras offer control over many of the options that physical cameras have. You can control focal length, lens distortion, f-stop, film back, and shutter speed. Although they are simple, cameras are the window into a scene, so using them properly will improve your art and add impact to your designs.

Navigating the modo User Interface

Modo offers visual cues that enable us to keep things straight from an orientation standpoint. In the bottom-left corner of each viewport window, a small axis indicator shows, in the orthographic view, the two axes making up the plane of view. The colored lines point in the positive direction, and the colors always correspond to a particular direction

(red for x, green for y, and blue for z). These same colors will appear in tool handles after modeling begins. In the perspective view, indicators are displayed for all three axes. As in the orthographic view, each line indicates the positive direction. In the perspective view, there is also a light gray square that aligns itself to the two axes that are most perpendicular to the current view. This square indicates the orientation of the work plane, which is presented in more depth in Chapter 2, "Creating Objects." By staying aware of these markers, we can more easily keep the scene aligned.

In addition to the axis widget, the perspective view offers a gradient background that helps to keep us from looking at our scene upside-down (or enables us to more easily get there if that is our desire). The background consists of a two-color gradient: a light blue-gray color indicates up, in the positive y direction (think of the sky); and a darker shade of the bluish color fills the negative y direction (indicating the ground). Because we spend a lot of time rotating around our objects to get the best view, staying aware of this gradient ensures that we keep our feet on the ground, so to speak.

Before you move on to navigating this space, you need to know about scale. You may notice that there are no document boundaries in 3D space as there are in a page layout document, at least not visible ones. What we do have to consider is the depth that the computer is able to re-create. This is called *draw distance*. A good example of a short draw distance can be seen in older racing video games. As you drive along a course and look into the distance, objects (such as buildings, trees, and mountains) will appear rather suddenly instead of growing from small points on the horizon. Although modern 3D applications are much more capable of handling distance than those games were, we want to keep draw distance in mind and create our scenes at (or near) actual scale.

In the bottom-right corner of each modo viewport, a display shows the scale of the small grid boxes. By default, the three orthographic views are linked together, but the perspective view is independent. Because modo uses physical scale for many aspects of lighting and texturing, it is important to check your scale as you begin to create models in 3D space. I have seen many students (and, regrettably, myself) create large sections of scenes only to realize that the pencil onscreen is as large as an oak tree—or bigger! In the 3D view, you will also see a light-colored grid that changes position and orientation based on your perspective. This is the Work Plane, and it is a huge help when you begin modeling objects.

Figure 1.6 is a breakdown of the modo user interface (UI) with labels for the features that are pertinent to this section.

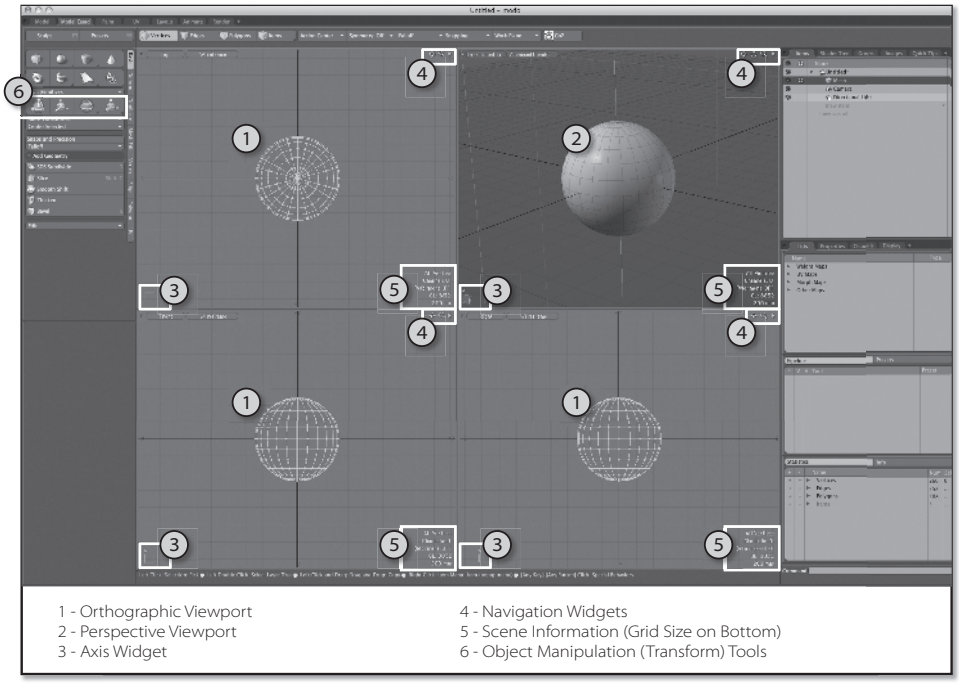
1 - Orthographic Viewport
2 - Perspective Viewport
3 - Axis Widget
4 - Navigation Widgets
5 - Scene Information (Grid Size on Bottom)
6 - Object Manipulation (Transform) Tools

Figure 1.6
The modo user interface

Maneuvering in 3D Space

Now that you have a basic understanding of our canvas, you are ready to start maneuvering the viewports to get the best view for any part of the creative process. In a 2D application, the view controls are simple and are limited to panning, zooming, and rotating the canvas. Interestingly, the addition of just one more navigational feature significantly complicates the way that we interact with the environment.

Panning and zooming are relative to the perspective of the user, so they do not change much. There are two ways to pan in 3D space: by clicking and dragging on the pan icon in the upper-right corner of the view or by holding Alt (Windows) or Option (Mac) while clicking and dragging.

Zooming can be achieved via one of three methods. As with the pan tools, there is a zoom widget in the upper-right corner (this one zooms directly toward the center of the view). Holding Alt+Ctrl (Windows) or Option+Control (Mac) while clicking and dragging zooms based on the position of the cursor. By using this method with the right mouse button, a zoom area can be selected. This creates a box, and the area inside will

zoom to fill the view after the mouse button is released. Finally, you can zoom with the scroll wheel. This last option is also context sensitive and will zoom toward the cursor position.

> While in the camera view, using the Alt+Ctrl / Option+Control key combo with the right mouse button adjusts the focal length of the camera and enables you to zoom in and out, as opposed to actually moving the camera in and out in space.

The real complication comes with rotation. Rotation works based on the combination of two axes. Consider the way we rotate in two dimensions. With the x and y axes covering the screen horizontally and vertically, we can rotate along the perpendicular axis (in this case, z). Because only two axes are present, we can rotate in only one dimension. The addition of the third axis adds two more possible rotation options, as we now have three planes to consider (xy, xz, and yz). As if this did not complicate things enough, there is one more point to consider: the center of rotation. In a 2D layout, we rotate relative to the center of the document. If we were to rotate only in relation to the origin (center of space), we would be very limited in our access to the work area.

To move freely in a 3D workspace, we need to rotate our viewport dynamically. So there are three types of rotation to deal with:

- Around the view focal point
- On a virtual tripod
- Rolling around the axis perpendicular to the perspective

Figure 1.7 shows a model from the front. Figure 1.8 shows that same object with the view rotated to show the depth of the model.

Figure 1.7
Flat front view of a model

Figure 1.8
Perspective view of a model

First, let's look at the rotation based on the center. This type of rotation is accessed by pressing the Alt/Option key while clicking the left mouse button, or by clicking and dragging the rotation widget in the top-right corner of the viewport. As you drag up, down, left, and right, the view shifts around the focal point. This rotation technique keeps a set distance from the center of view and always faces that point in space until another focus is selected.

To rotate the view from a tripod, press Alt/Option and the right mouse button. The tripod rotation works exactly like a real tripod: the orientation of the view is changed, but the position of the view remains constant. This can be quite useful when working on large scenes or architectural interiors. One thing to remember is that this option is specific to views from cameras and lights, so the basic perspective view is not able to use tripod-based rotation.

QUICK TURNTABLE

Using the right mouse button to rotate in the perspective view enables the view to rotate on its own. As soon as you release the button, the view continues to spin, with the speed based on the mouse speed when the button was released. Moving slowly allows for slow and subtle rotation, while a quick flick of the mouse sends the view spinning rapidly. There is a falloff of speed, and then the view comes to rest. This can also be done with the question mark (?) key, which gives a single revolution around the scene.

Finally, by holding Alt/Option and using the middle mouse button (or scroll wheel button), the view can spin perpendicular to the viewport. This can be useful if the current view is upside-down and needs to be flipped over quickly. This is also useful when aligning a view that is slightly skewed and requires a minor rotation to see a level view of the scene.

The three movement tools work well when the view is focused on either your entire scene or a specific selection. The A key is used to center on all visible items in the scene. You can use the rotation tool to spin around your entire scene, and press Shift+A to focus on a selection. By selecting the area you are interested in or working on, you can center your perspective and rotate around the area in question. Subsequent chapters cover selections in more depth, but you can get started selecting by simply clicking and dragging across some surfaces (polygons, edges, or points) in your scene. After you select something, pressing Shift+A will center your view on that selection.

Practice: Navigating in Space

Open the file `Navigation_Practice.lxo` from the included DVD. Spend a few minutes moving around the scene. The more you navigate the space, the more comfortable it will become. Start by moving around the objects generally, and then choose various sections of the objects and manipulate the perspective until you get a good view of them. See if you can get a side view showing all three objects aligned in the view. Remember to utilize all of your newly learned navigation tools. Rotate, pan, and zoom your view. Use the A key and Shift+A centering tools to center your view on an area of interest.

The practice file starts out with the viewport containing all three objects from an angle, as seen in Figure 1.9. Try to duplicate the view shown in Figure 1.10. A few minutes practicing with a simple scene such as this will reduce frustration when you have a more complex scene and are still familiarizing yourself with the controls.

Moving Objects in 3D Space

Now that we have discussed moving around, you can begin to look at moving the objects that you create. There are three methods for basic object manipulation: move, scale, and rotate. Each of these tools (known as *transform tools*) can be activated by clicking the corresponding button on the left side, toward the top of the user interface, or by using their hot keys. Each transform tool has some quirks or additional features that will speed your workflow if you take advantage of them.

After a transform tool is activated, a property tab appears in the bottom-left corner. The fields in this tab enable you to enter numeric values for each of the transform functions as well as control of some additional options for each tool. These values can be entered by clicking in the fields and entering a value, by clicking on the arrows to increase or decrease the values by small increments, or by clicking and dragging the arrows.

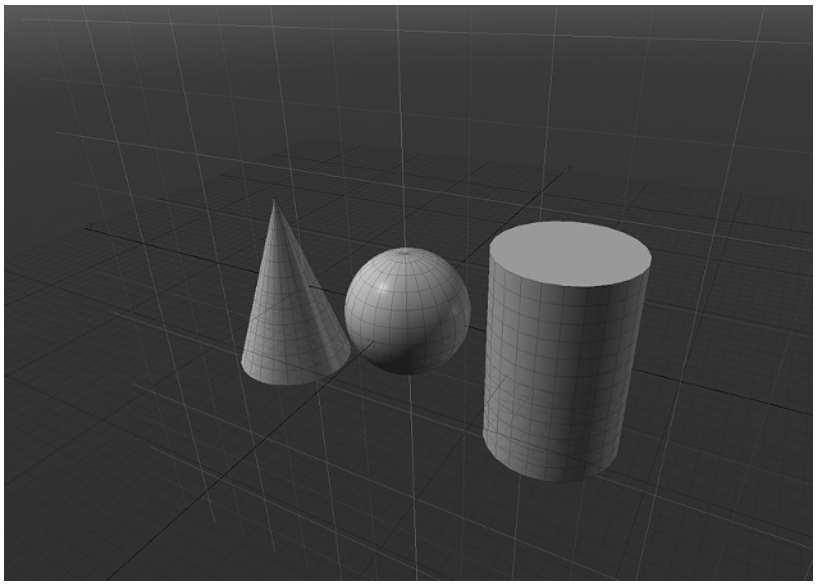

Figure 1.9
Starting perspective

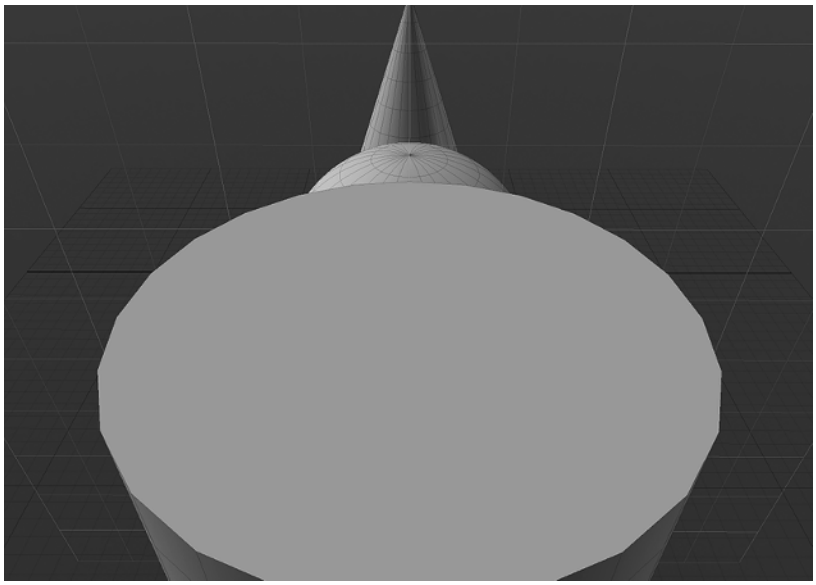

Figure 1.10
Goal after navigating the perspective view

If a value is entered in a numeric field, holding the Ctrl / Control key and pressing the Enter key (on the Mac, Tab or Enter) will change the other values proportionally. In other words, if all fields read 100% for scale and you enter 10% in the X value, the other two values will be changed to 10% automatically.

First, let's consider the Move tool. When activated, the tool handles appear at the center of the selection and show arrows that align to each axis (x, y, and z) and circles that float between the arrows in free space. For reference, the colors of the arrows correspond to the axis of movement. To move in a single axis, click and drag on the individual arrows. When the axis is active, the arrow turns yellow. Clicking and dragging anywhere but on the tool handles in the perspective view causes the handles to snap to the click point, and dragging moves the object on the Work Plane. Clicking and dragging on the circles moves the object perpendicular to the axis of the same color (red, green, or blue for x, y, and z, respectively). While in the orthographic views, clicking and dragging on the handles produces the same behavior as in the perspective view, and clicking off the handles moves the object along the plane in view. The hot key for the Move tool is W.

CENTERING THE TOOL HANDLES

The default behavior of the tool handles to align to a click point is controlled by the Action Center, which is set to Automatic when modo starts up. I cover Action Centers in depth in future chapters, but for now, if the behavior is difficult to use, you can click Action Center above the viewports and select Selection Center Auto Axis. This ensures that the tool handles remain in the center of the object and still align to the x, y, and z axes.

Next, let's consider the Rotate tool. The handles for this tool are circles. Each circle is colored like the circles on the Move tool. Clicking and dragging on these rotates the object perpendicular to the axis of the corresponding color. Again, the active handle turns yellow. Clicking off the handles snaps the tool to the click point, and dragging free-rotates the object in all directions at once. This type of rotation is very difficult to control and is not recommended. The gray circle that encompasses the rest of the handles rotates the view and rotates the object perpendicular to the current view. In the orthographic viewports, the gray circle is no different from the colored handles, but in the perspective views, this will rotate variably based on the angle of view. This is a pretty special use, but when it is needed, it can come in quite handy! Because rotation by its very nature takes place around an axis, there is no way to rotate in a single direction, and so there are no separate axes at one time (as with the Move tool). The hot key for the Rotate tool is E.

Pressing and holding the Ctrl / Control key prior to clicking a rotation handle causes the angle of rotation to snap to 15-degree increments. This is useful when precise rotation is required.

The Scale tool changes the size of the selected object or objects. The handles for the Scale tool are similar to those of the Move tool. The Scale tool is visually differentiated

by the ends of the tool handles, which are boxes instead of arrows. Just as on the Move tool, the handles scale in one direction, and the circles scale in two directions (one plane). Clicking off the handles scales based on the work plane. Because clicking in open space still scales the object independently, it is usually preferable to use the planar circles and scale uniformly in two directions. Unlike the Move tool, the Scale tool has another behavior controlled by the cyan-colored circle at its center. This circle scales the object uniformly in all directions. The hot key for the Scale tool is R.

The Transform tool is a combination of all three of the other transform tools (Move, Scale, and Rotate). The question you may ask is, "Why on earth would I use individual tools when the Transform tool does it all?" The answer is that it really doesn't do it all. Although the tool does provide the basic function of the Move, Rotate, and Scale tools, it lacks several key options. The move portion of the Transform tool does not have planar handles. The scale portion lacks both planar handles and the uniform scale option. Because the tool has to act on one axis at a time, it is often more productive to switch tools and be able to scale or move in multiple (or all) directions at once. The Transform tool comes in most handy when making quick adjustments to both movement and rotation. Mostly, this comes down to a matter of personal preference. Try the tools and see which ones make the most sense to you and allow you to work the most efficiently. The hot key for the Transform tool is Y.

Tools in modo are "sticky" and will remain active until the tool is dropped. A tool can be dropped by pressing the spacebar, the Q key, or the Esc key. On subsequent presses, the spacebar will switch modes between vertex, edge, and polygon. The Esc key will (with additional strokes) clear out the tool pipe. I cover these functions in the next few chapters, but because the Q key is bound only to dropping a tool, it is often the first choice for this function.

Transform tools applied in vertex, edge, or polygon mode will alter the position of geometry and cannot be reset with the exception of centering the selection by using the Center Selected tool (under the transform tools on the left side of the screen). Objects can be centered on any axis or combination of axes. When transforming objects in Item mode, the changes are logged under the properties for the mesh layer. The Properties tab at the bottom-right corner of the screen contains numeric fields for position, scale, and rotation. After a change has been made, it can be adjusted or reset in the numeric fields on this tab.

Practice: Moving Objects

Open the file `Transform_Practice.lxo` from the included DVD. The file contains six cubes in individual layers. When in individual layers, each object can be moved separately when in Item mode.

1. Enter Item mode by clicking the Items button above the viewports (next to Vertices, Edges, and Polygons).

2. Select a layer either by clicking on it in a viewport or by clicking on it in the Items list tab on the upper-right side of the screen.
3. To move more than one layer at a time, Shift+click on it in the item list.
4. Use combinations of the Move, Rotate, Scale, and Transform tools to get the cubes into the positions shown in Figure 1.11.

Figure 1.11
Final position and orientation of the six cubes

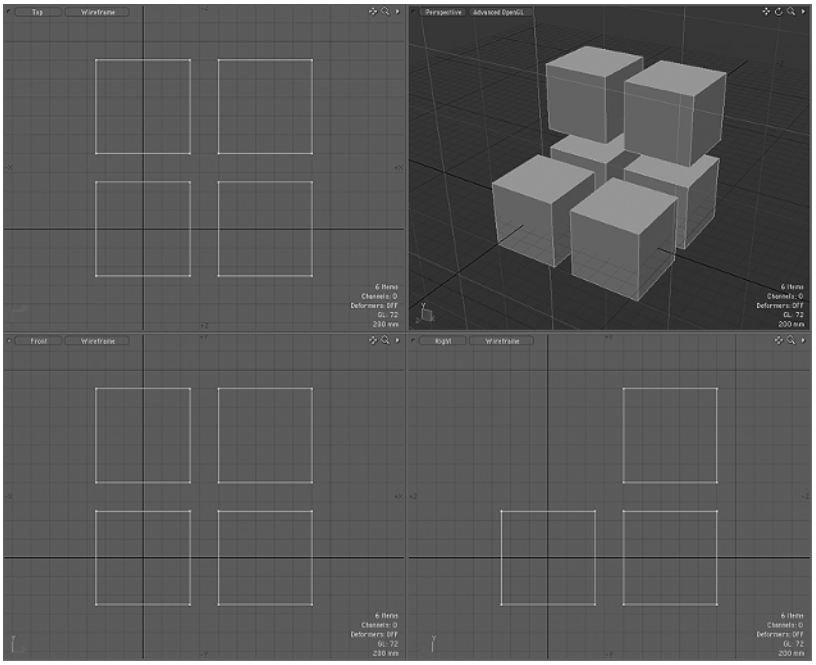

Review

This chapter has covered an intro to 3D space as well as the basics of moving around our viewports and objects in three dimensions. You have looked at some of the basic sections of the user interface (UI) and learned some important hot keys. With the large number of tabs, menus, and buttons in the interface, learning hot keys can be important to a quick and efficient workflow. Here are the hot keys, to recap:

- Alt+Shift+click / Option+Shift+click = pan view
- Alt+Ctrl+click / Option+Control+click = zoom view
- Alt+click / Option+click = rotate view
- A = center view on all visible items
- Shift+A = center view on selection

- W = Move tool
- E = Rotate tool
- R = Scale tool
- Y = Transform tool

In the next chapter, you will start creating 3D objects. You will look at object primitives and all of their options that enable you to control their size proportions and structure. This will also give you the opportunity to explore some additional sections of the modo user interface.

On the DVD for this chapter are practice files and videos covering the topics discussed in the previous pages. These short videos show tools and procedures in action to help accelerate the learning process.